CSE 332

Data Abstractions

More Parallel Primitives and Parallel Sorting

Outline

Maps and Reductions   1

Reductions

INPUT: An array
OUTPUT: A combination of the array by an associative operation
The general name for this type of problem is a **reduction**. Examples
include: max, min, has-a, first, count, sorted

Maps

INPUT: An array
OUTPUT: Apply a function to every element of that array
The general name for this type of problem is a **map**. You can do this
with any function, because the array elements are independent.

Today, we'll add in two more:

- Scan
- Pack (or filter)

As we'll see, both of these are quite a bit less intuitive **in parallel** than
map and reduce.

Scan

Suppose we have an associative operation $\oplus$ and an array a:

$$a: \quad \boxed{a_0 \mid a_1 \mid a_2 \mid a_3}$$
a[0] a[1] a[2] a[3]

Then, `scan(a)` returns an array of "partial sums" (using $\oplus$):

$$\text{scan(a):} \quad \boxed{a_0 \mid a_0 \oplus a_1 \mid a_0 \oplus a_1 \oplus a_2 \mid a_0 \oplus a_1 \oplus a_2 \oplus a_3}$$
b[0] b[1] b[2] b[3]

It's hard to see at first, but this is actually a really powerful tool. It gives us a "partial trace" of the operation as we apply it to the array (for free).

No Seriously

splitting, load balancing, quicksort, line drawing, radix sort, designing binary adders, polynomial interpolation, decoding gray codes

---

For the sake of being clear, we'll discuss scan with $\oplus = +$.
That is, "prefix sums" of an array":

Example (Prefix Sum)

$$a: \quad \boxed{5 \mid 1 \mid 3 \mid 4 \mid 2}$$
a[0] a[1] a[2] a[3] a[4]

$$\text{scan(a):} \quad \boxed{5 \mid 6 \mid 9 \mid 13 \mid 15}$$
b[0] b[1] b[2] b[3] b[4]

Sequential Code

```
1  int[] prefixSum(int[] input) {
2      int[] output = new int[input.length];
3      int sum = 0;
4      for (int i = 0; i < input.length; i++) {
5          sum += input[i];
6          output[i] = sum;
7      }
8      return output;
9  }
```

If you have a really good memory, you'll remember that on the **very first day of lecture**, we discussed a very similar problem.

Sequential Code

```
1  int[] prefixSum(int[] input) {
2      int[] output = new int[input.length];
3      int sum = 0;
4      for (int i = 0; i < input.length; i++) {
5          sum += input[i];
6          output[i] = sum;
7      }
8      return output;
9  }
```

Bad News

This **algorithm** does not parallelize well. Step i needs the outputs from all the previous steps. This might as well be an algorithm on a linked list.
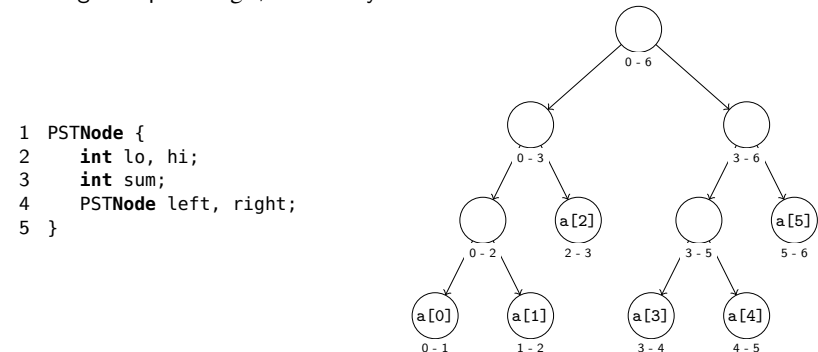
So, what do we do?

Come Up With A Better Algorithm!

The solution here will be to add a "pre-processing step". This is essentially what we did in the first lecture.

We begin with an array as usual:

$$a: \quad \boxed{ \mid \mid \mid \mid \mid }$$
a[0] a[1] a[2] a[3] a[4] a[5]

Then, transform it into a **balanced tree**, because $\lg n$ height will allow us to get a span of $\lg n$, eventually:

```
1  PSTNode {
2      int lo, hi;
3      int sum;
4      PSTNode left, right;
5  }
```

Creating the tree is a standard divide-and-conquer recursive algorithm:

```
1 PSTNode {
2     int lo, hi;
3     int sum;
4     PSTNode left, right;
5 }
```
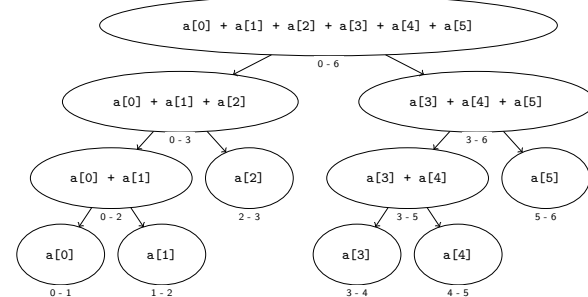


```
1  PSTNode processInput(int[] input, int lo, int hi) {
2      if (hi - lo == 1) {
3          return new PSTNode(lo, hi, input[lo]);
4      }
5      else {
6          mid = lo + (hi - lo)/2;
7          PSTNode left = processInput(lo, mid);
8          PSTNode right = processInput(mid, hi);
9          return new PSTNode(lo, hi, left.sum + right.sum, left, right);
10     }
11 }
```
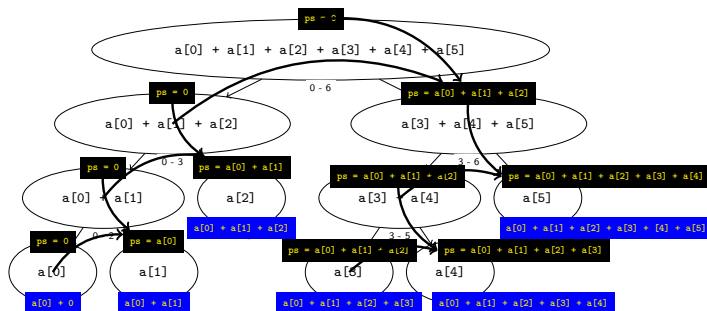
Now, we have the entire tree filled out:

To fill in all the prefix sums, we recursively fill them in down the tree.
Since the non-leaf nodes don't have access to the elements of the array,
we fill in a **pre-scan** (everything up to, but not including the range).

To fill in all the **pre-scans**, we recursively fill them in down the tree:


```
1  void makeOutput(int[] output, PSTNode current, int prescan) {
2      if (current is a leaf) {
3          output[current.lo] = prescan + current.sum;
4      }
5      else {
6          makeOutput(output, current.left, prescan);
7          makeOutput(output, current.right, prescan + current.left.sum);
8      }
9  }
```

Adding a sequential cut-off isn't too bad:

Processing the Input

This is just a normal sequential cut-off. The leaves end up being cutoff size ranges instead of ranges of one.

Constructing the Output

We must sequentially compute the prefix sum at our leaves as well:

```
1  output[lo] = prescan + input[lo];
2  for (i = lo + 1; i < hi; i++) {
3      output[i] = output[i-1] + input[i]
4  }
```

Notice that this means we must pass the input array to this phase now.

Here the idea is that we'd like to filter the array given some predicate (e.g., ≤ 7). More specifically:

Pack/Filter

Suppose we have a function $f : E \rightarrow$ boolean and an array a of type E:

a: | a_0 | a_1 | a_2 | a_3 |
|---|---|---|---|
a[0] | a[1] | a[2] | a[3]

Then, pack(a) returns an array of elements x for which $f(x) =$ true. For example, if arr = [1, 3, 8, 6, 7, 2, 4, 9] and f(x) = x % 2 == 0, then pack(arr) = [8, 6, 2, 4].

The key to doing this in parallel is scan!

Let f(x) = x % 2 == 0.

Parallel Pack

input: | 1 | 3 | 8 | 6 | 7 | 2 | 4 | 9 |
|---|---|---|---|---|---|---|---|
a[0] | a[1] | a[2] | a[3] | a[4] | a[5] | a[6] | a[7]

1 Use a **map** to compute a bitset for $f(x)$ applied to each element

bitset: | 0 | 0 | 1 | 1 | 0 | 1 | 1 | 0 |
|---|---|---|---|---|---|---|---|
b[0] | b[1] | b[2] | b[3] | b[4] | b[5] | b[6] | b[7]

2 Do a **scan on the bit vector** with $\oplus = +$:

bitsum: | 0 | 0 | 1 | 2 | 2 | 3 | 4 | 4 |
|---|---|---|---|---|---|---|---|
c[0] | c[1] | c[2] | c[3] | c[4] | c[5] | c[6] | c[7]

3 Do a **map on the bit sum** to produce the output:

output: | 8 | 6 | 2 | 4 |
|---|---|---|---|
d[0] | d[1] | d[2] | d[3]

```
1 output = new E[bitsum[n-1]];
2 for (i=0; i < input.length; i++) {
3     if (bitset[i] == 1) {
4         output[bitsum[i] - 1] = input[i];
5     }
6 }
```

- We can combine the first two passes into one (just use a different base case for prefix sum)

- We can also combine the third step into the second part of prefix sum

- Overall: $\mathcal{O}(n)$ work and $\mathcal{O}(\lg n)$ span. (Why?)

We can use scan and pack in all kinds of situations!

```
1 int[] quicksort(int[] arr) {
2     int pivot = choosePivot();
3     int[] left = filterLessThan(arr, pivot);
4     int[] right = filterGreaterThan(arr, pivot);
5     return quicksort(left) + quicksort(right);
6 }
```

Do The Recursive Calls in Parallel

Assuming a good pivot, we have:

$$\text{work}(n) = \begin{cases} \mathcal{O}(1) & \text{if } n = 1 \\ 2\text{work}(n/2) + \mathcal{O}(n) & \text{otherwise} \end{cases}$$

and

$$\text{span}(n) = \begin{cases} \mathcal{O}(1) & \text{if } n = 1 \\ \max(\text{span}(n/2), \text{span}(n/2)) + \mathcal{O}(n) & \text{otherwise} \end{cases}$$

These solve to $\mathcal{O}(n \lg n)$ and $\mathcal{O}(n)$. So, the parallelism is $\mathcal{O}(\lg n)$.

```
1  int[] quicksort(int[] arr) {
2    int pivot = choosePivot();
3    int[] left = filterLessThan(arr, pivot);
4    int[] right = filterGreaterThan(arr, pivot);
5    return quicksort(left) + quicksort(right);
6  }
```

Do The Partition in Parallel

The partition step is just two filters or packs. Each pack is $\mathcal{O}(n)$ work, but $\mathcal{O}(\lg n)$ span! So, our new span recurrence is:

$$\text{span}(n) = \begin{cases} \mathcal{O}(1) & \text{if } n = 1 \\ \max(\text{span}(n/2), \text{span}(n/2)) + \mathcal{O}(\lg n) & \text{otherwise} \end{cases}$$

Master Theorem says this is $\mathcal{O}(\lg^2 n)$ which is neat!

```
1  int[] mergesort(int[] arr) {
2    int[] left = getLeftHalf();
3    int[] right = getRightHalf();
4    return merge(mergesort(left), mergesort(right));
5  }
```

Do The Recursive Calls in Parallel

This will get us the same work and span we got for quicksort when we did this:

- work$(n) = \mathcal{O}(n \lg n)$
- span$(n) = \mathcal{O}(n)$
- Parallelism is $\mathcal{O}(\lg n)$

Now, let's try to parallelize the merge part.

As always, when we want to parallelize something, we can turn it into a divide-and-conquer algorithm.

### Do The Merge in Parallel

Merge takes as input two arrays:

arr[0]	arr[1]	arr[2]	arr[3]	arr[4]

arr[5]	arr[6]	arr[7]	arr[8]

1. Find the median of the **larger** array (just the middle index):

		X		
arr[0]	arr[1]	arr[2]	arr[3]	arr[4]

arr[5]	arr[6]	arr[7]	arr[8]

2. Partition the **smaller** array using X as a pivot. To do this, **binary search** the smaller array:

		X		
arr[0]	arr[1]	arr[2]	arr[3]	arr[4]

		Y	
arr[5]	arr[6]	arr[7]	arr[8]

3. Now, we have four pieces $\leq X$, $> X$, $\leq Y$, and $> Y$. In the sorted array, the $\leq$ pieces will be entirely before the $>$ pieces.

		X			Y
arr[0]	arr[1]	arr[2]	arr[5]	arr[6]	arr[7]

arr[3]	arr[4]	arr[8]

4. Recursively apply the merge algorithm (until some cut-off)!

First, we analyze **just the parallel merge**:

Parallel Merge Analysis

The non-recursive work is $\mathcal{O}(1) + \mathcal{O}(\lg n)$ to find the splits.

The **worst case** is when we split the bigger array in half and the smaller array is all on the left (or all on the right). In other words:

$$\text{work}(n) \leq \begin{cases} \mathcal{O}(1) & \text{if } n = 1 \\ \text{work}(3n/4) + \text{work}(n/4) + \mathcal{O}(\lg n) & \text{otherwise} \end{cases}$$

and

$$\text{span}(n) \leq \begin{cases} \mathcal{O}(1) & \text{if } n = 1 \\ \max(\text{span}(3n/4) + \text{span}(n/4)) + \mathcal{O}(\lg n) & \text{otherwise} \end{cases}$$

These solve to work$(n) = \mathcal{O}(n)$ and span$(n) = \mathcal{O}(\lg^2 n)$.

Now, we calculate the work and span of **the entire parallel mergesort**.

Putting It Together

$$\text{work}(n) = \mathcal{O}(n \lg n)$$

$$\text{span}(n) \leq \begin{cases} \mathcal{O}(1) & \text{if } n = 1 \\ \text{span}(n/2) + \mathcal{O}(\lg^2 n) & \text{otherwise} \end{cases}$$

This works out to $\text{span}(n) = \mathcal{O}(\lg^3 n)$.

This isn't quite as much parallelism as quicksort, but **this one is a worst case guarantee**!